

Classical Automata on Promise Problems*

Viliam Geffert[†]

Department of Computer Science, P. J. Šafárik University, Košice, Slovakia
viliam.geffert@upjs.sk

Abuzer Yakaryılmaz[‡]

National Laboratory for Scientific Computing, Petrópolis, RJ, Brazil
abuzer@lncc.br

October 17, 2014

Abstract

Promise problems were mainly studied in quantum automata theory. Here we focus on state complexity of classical automata for promise problems. First, it was known that there is a family of unary promise problems solvable by quantum automata by using a single qubit, but the number of states required by corresponding one-way deterministic automata cannot be bounded by a constant. For this family, we show that even two-way nondeterminism does not help to save a single state. By comparing this with the corresponding state complexity of alternating machines, we then get a tight exponential gap between two-way nondeterministic and one-way alternating automata solving unary promise problems. Second, despite of the existing quadratic gap between Las Vegas realtime probabilistic automata and one-way deterministic automata for language recognition, we show that, by turning to promise problems, the tight gap becomes exponential. Last, we show that the situation is different for one-way probabilistic automata with two-sided bounded-error. We present a family of unary promise problems that is very easy for these machines; solvable with only two states, but the number of states in two-way alternating or any simpler automata is not limited by a constant. Moreover, we show that one-way bounded-error probabilistic automata can solve promise problems not solvable at all by any other classical model.

Keywords: descriptonal complexity, promise problems, nondeterministic automata, probabilistic automata, alternating automata

1 Introduction

Promise problem is a generalization of language recognition. Instead of testing all strings over a given alphabet for language membership, we focus only on a given subset of strings as potential inputs that should be tested. The language under consideration and its complement form this subset.

Promise problems have already provided several different perspectives in computational complexity. (See the survey by Goldreich [16].) For example, it is not known whether the class of languages recognizable by bounded-error probabilistic polynomial-time algorithms has

*A preliminary version of this work was presented at the 16th International Workshop on Descriptonal Complexity of Formal Systems (DCFS 2014), August 5–8, 2014, Turku, Finland [vol. 8614 of LNCS, pp. 126–137, Springer-Verlag, 2014]. An ArXiv version is at <http://arxiv.org/abs/1405.6671>.

[†]Partially supported by the Slovak grant contracts VEGA 1/0479/12 and APVV-0035-10.

[‡]Partially supported by CAPES with grant 88881.030338/2013-01, ERC Advanced Grant MQC, and FP7 FET project QALGO. Moreover, the part of the research work was done while Yakaryılmaz was visiting Kazan Federal University.

a complete problem, but the class of promise problems solvable by the same type of algorithms has some complete problems. A similar phenomenon has been observed for quantum complexity classes [37]. The first known result on promise problems for restricted computational models we are aware of was given by Condon and Lipton in 1989 [7]. They showed that a promise version of emptiness problem is undecidable for one-way probabilistic finite automata. In the literature, some separation results regarding restricted computational models have also been given in the form of promise problems. (See, e.g., [10, 9].) The first result for restricted quantum models was given by Murakami et al. [29]: There exists a promise problem that can be solved by quantum pushdown automata exactly but not by any deterministic pushdown automaton. Recently, Ambainis and Yakaryılmaz [2] showed that there is an infinite family of promise problems which can be solved exactly just by tuning transition amplitudes of a realtime two-state quantum finite automaton (QFA), whereas the size of the corresponding classical automata grows above any fixed limit. Moreover, Rashid and Yakaryılmaz [33] presented many superiority results of QFAs over their classical counterparts. There are also some new results on succinctness of QFAs [40, 17, 18, 39, 3] and a new result on quantum pushdown automata [30].

In this paper, we turn our attention to classical models and obtain some new results. In the next two sections, we give some preliminaries and present some basic facts for the classical models. Among others, we show that (i) the computational power of deterministic, nondeterministic, alternating, and Las Vegas probabilistic automata is the same, even after turning from the classical language recognition to solving promise problems, and, (ii) on promise problems, neither existing gaps of succinctness for classical language recognition can be violated, between any two of deterministic, nondeterministic, and alternating automata models.

Then, in Section 4, we consider a family of unary promise problems given by Ambainis and Yakaryılmaz in [2], solvable by only two-state realtime quantum finite automata. On the other hand, for solving this family by the use classical automata, we show that the exact number of states for one-way/two-way deterministic/nondeterministic automata is the same. Thus, for this family, replacing the classical one-way deterministic automata even by two-way nondeterministic automata does not help to save a single state. However, for the same family, already the one-way alternation does help, which gives us the tight exponential gap between one-way alternating and two-way sequential models of automata solving unary promise problems.

In Section 5, we show that the trade-off for the case of Las Vegas probabilistic versus one-way deterministic automata is different; by turning from language recognition to promise problems the tight quadratic gap changes to a tight exponential gap.

In Section 6, we show how the situation changes for finite state automata with two-sided bounded-error. First, we present a probabilistically very easy family of unary promise problems, i.e., solvable by one-way two-state automata with bounded-error. Then, we show that this family is hard for two-way alternating automata, and hence for any simpler classical automata as well. Finally, we prove that, on promise problems, bounded-error probabilistic automata are more powerful than any other classical model.

2 Preliminaries

We assume the reader is familiar with the basic standard models of finite state automata (see e.g. [19]). For a more detailed exposition related to probabilistic automata, the reader is referred to [31, 4]. Here we only recall some models and introduce some elementary notation.

By Σ^* , we denote the set of strings over an alphabet Σ . The set of strings of length n is denoted by Σ^n and the unique string of length zero by ε . By \mathcal{N}_+ , we denote the set of all positive integers. The cardinality of a finite set S is denoted by $\|S\|$.

A *one-way nondeterministic finite automaton* with ε -moves (1_ε NFA, for short) is a quintuple $\mathcal{A} = (S, \Sigma, H, s_1, S_A)$, where S is a finite set of states, Σ an input alphabet, $s_1 \in S$ an initial

state, and $S_A \subseteq S$ a set of accepting (final) states. Finally, $H \subseteq S \times (\Sigma \cup \{\varepsilon\}) \times S$ is a set of *transitions*, with the following meaning. $s \xrightarrow{a} s' \in H$: if the next input symbol is a , \mathcal{A} gets from the state s to the state s' by reading a from the input. $s \xrightarrow{\varepsilon} s' \in H$: \mathcal{A} gets from s to s' without reading any symbol from the input.¹

The machine \mathcal{A} *halts* in the state s , if there are no executable transitions in H at a time. Typically, this happens after reading the entire input, but the computation may also be blocked prematurely, by *undefined transitions*, i.e., by absence of transitions for the current state s and the next input symbol a (or ε).

An *accepting computation* starts in the initial state s_1 and, after reading the entire input, it halts in any accepting state $s' \in S_A$. The language *recognized* (accepted) by \mathcal{A} , denoted by $L(\mathcal{A})$, is the set of all input strings having at least one accepting computation. The inputs with no accepting computation are *rejected*.

A *two-way nondeterministic finite automaton* (2NFA) is defined in the same way as 1_ε NFA, but now \mathcal{A} can move in both directions along the input. For these reasons, the transitions in H are upgraded as follows. $s \xrightarrow{a\rightarrow} s'$, $s \xrightarrow{a\leftarrow} s'$, $s \xrightarrow{a\circ} s'$: if the current input symbol is a , \mathcal{A} gets from s to s' and moves its input head one position to the right, to the left, or keeps it stationary, respectively. The input is enclosed in between two special symbols “ \vdash ” and “ \dashv ”, called the *left and right endmarkers*, respectively. By definition, the input head cannot leave this area, i.e., there are no transitions moving to the left of \vdash nor to the right of \dashv . \mathcal{A} starts in s_1 with the input head at the left endmarker and accepts by halting in $s' \in S_A$ anywhere along the input tape.

A *deterministic finite automaton* (1_ε DFA or 2DFA) can be obtained from 1_ε NFA or 2NFA by claiming that the transition set H does not allow executing more than one possible transition at a time. Consequently, a 1_ε DFA \mathcal{A} can have at most one transition $s \xrightarrow{a} s'$, for each $s \in S$ and each $a \in \Sigma$. In theory, one can consider ε -transitions (even though this feature is rarely utilized): a transition $s \xrightarrow{\varepsilon} s'$ implies that there are no other transitions starting in the same state s . Similar restrictions can be formulated for 2DFAs.

A *self-verifying finite automaton* (1_ε SVFA or 2SVFA) [22] is a 1_ε NFA or 2NFA \mathcal{A} which is additionally equipped with $S_R \subseteq S$, the set of rejecting states disjoint from S_A , the set of accepting states. The remaining states form $S_N = S \setminus (S_A \cup S_R)$, the set of neutral “don’t know” states. In this model, (i) for each accepted input, \mathcal{A} has at least one computation path halting in an accepting state and, moreover, no path halts in a rejecting state, (ii) for each rejected input, \mathcal{A} has at least one computation path halting in a rejecting state and, moreover, no path halts in an accepting state. In both these cases, some computation paths may return “don’t know” answers, by halting in neutral states or by getting into infinite loops.

A *one-way probabilistic finite automaton* with ε -moves (1_ε PFA) is defined in the same way as 1_ε NFA, but the transitions in H are upgraded with probabilities, as follows. $s \xrightarrow{a,p} s' \in H$: if the current input symbol is a , \mathcal{A} gets from the state s to the state s' by reading this symbol with probability $p \in [0, 1]$. For $a = \varepsilon$, no input symbol is consumed. A transition with $p = 1$ may be displayed as $s \xrightarrow{a} s'$, since the next step is deterministic; transitions with $p = 0$ can be completely erased, as not executable. By definition,² for each $s \in S$ and each $a \in \Sigma$, the sum of probabilities over all transitions beginning in s and labeled by $a' \in \{a, \varepsilon\}$ must be equal to 1. The overall accepting probability of \mathcal{A} on a given input $w \in \Sigma^*$, denoted by $f_{\mathcal{A}}(w)$, is calculated over all accepting paths. Hence, the overall rejecting probability is $1 - f_{\mathcal{A}}(w)$.

A two-way version, 2PFA, was first introduced in [25].

¹Traditionally [19], the behavior is defined by a function $\delta : S \times (\Sigma \cup \{\varepsilon\}) \rightarrow 2^S$. However, a traditional transition $\delta(s, a) \ni s'$ can be easily converted to $s \xrightarrow{a} s'$ and vice versa. In the same way, all models introduced here are consistent with the traditional definitions. Using transitions instead of a δ -function, we just make the notation more readable.

²Traditionally [32], the behavior of PFAs is defined by stochastic matrices.

A *Las Vegas probabilistic finite automaton* (Las Vegas 1_ε PFA or Las Vegas 2PFA) [21] is a 1_ε PFA or 2PFA of special kind, obtained from self-verifying automata (i.e., from nondeterministic automata of special kind). Thus, transitions are upgraded with probabilities, and the state set is partitioned into the sets of accepting, rejecting, and neutral “don’t know” states. This gives three overall probabilities on a given input; accepting, rejecting, and neutral.

A *one-way alternating finite automaton* with ε -moves (1_ε AFA) is obtained from 1_ε NFA by partitioning the state set S into two disjoint subsets S_\exists and S_\forall , the sets of *existential* and *universal* states, respectively. The global rules are defined as is usual for alternating devices (see e.g. [5]): if, at the given moment, there are several executable transitions, the machine \mathcal{A} (i) nondeterministically chooses one of them, if it is in an existential state s , but (ii) follows, in parallel, all possible branches, if the current state s is universal.

For better readability, a nondeterministic switch from the existential s into one of the states s_1, \dots, s_k , by reading an input symbol a , will be displayed as $s \xrightarrow{a} s_1 \vee \dots \vee s_k$, while a parallel branching from the universal s to all these states as $s \xrightarrow{a} s_1 \wedge \dots \wedge s_k$. The same applies for $a = \varepsilon$.

The input is accepted, if all branches in the nondeterministically chosen computation subtree, rooted in the initial state at the beginning of the input and embedded in the full tree of all possible computation paths, halt in accepting states at the end of the input.³

Also this model was extended to a two-way version, 2AFA (see e.g. [13]).

A *realtime automaton* (rtDFA, rtSVFA, rtNFA, rtPFA, rtAFA, ...) is the corresponding one-way finite automaton⁴ that executes at most a constant number of ε -transitions in a row. (After that, the next input symbol is consumed or the machine halts.)

A *one-way ε -free automaton* (1DFA, 1SVFA, 1NFA, 1PFA, 1AFA, ...) is the corresponding one-way finite automaton with no ε -transitions at all. (This is a special case of a realtime machine, with the number of executed ε -transitions bounded by zero.)

It should be pointed out that, in 1_ε NFAs, the ε -transitions can be removed *without increasing* the number of states [26, Sect. 2.11]. Therefore, 1NFAs, rtNFAs, and 1_ε NFAs agree in upper/lower bounds for states. The same works for 1SVFAs, rtSVFAs, and 1_ε SVFAs, as well as for 1DFAs, rtDFAs, and 1_ε DFAs. For this reason, we consider only 1NFAs, 1SVFAs, and 1DFAs.

Occasionally, we shall also mention some restricted versions of two-way machines, e.g., *sweeping automata*, changing the direction of the input head movement only at the endmarkers [35, 24], or a very restricted version of sweeping automaton called a *restarting one-way automaton*, running one-way in an infinite loop [38]—if the computation does not halt at the right endmarker, then it jumps back to the initial state at the beginning of the input.

A *promise problem* is a pair of languages $P = (P_{\text{yes}}, P_{\text{no}})$, where $P_{\text{yes}}, P_{\text{no}} \subseteq \Sigma^*$, for some Σ , and $P_{\text{yes}} \cap P_{\text{no}} = \emptyset$ [37].

The promise problem P is *solved* by a finite automaton \mathcal{P} , if \mathcal{P} accepts each $w \in P_{\text{yes}}$ and rejects each $w \in P_{\text{no}}$. (That is, we do not have to worry about the outcome on inputs belonging neither to P_{yes} nor to P_{no} .) If $P_{\text{yes}} \cup P_{\text{no}} = \Sigma^*$ and P is solved by \mathcal{P} , we have a classical language recognition, and say that P_{yes} is *recognized* by \mathcal{P} .

If \mathcal{P} is a probabilistic finite automaton, we say that \mathcal{P} solves P with *error bound* $\varepsilon \in [0, \frac{1}{2}]$, if \mathcal{P} accepts each $w \in P_{\text{yes}}$ with probability at least $1 - \varepsilon$ and rejects each $w \in P_{\text{no}}$ with probability at least $1 - \varepsilon$. (The remaining probability goes to erroneous answers.) \mathcal{P} solves P with a

³The alternating automaton should not be confused with a *Boolean finite automaton*, which is (quite inappropriately) also sometimes called “alternating” in some literature. Instead of alternating existential and universal branching, the Boolean automaton can control acceptance by the use of arbitrarily complex Boolean functions. As an example, $s \xrightarrow{a} s_1 \wedge (s_2 \vee s_3)$ specifies that, for the current input symbol a , the subtree rooted in the state s is successful if and only if the subtree rooted in s_1 and at least one of those rooted in s_2, s_3 are successful.

⁴All realtime *quantum* models use a classical input head and so they are not allowed to create a superposition of the head positions on the input tape.

bounded-error, if, for some ε , it solves P with the error bound ε . If $\varepsilon = 0$, the problem is solved *exactly*.

A special case of bounded-error is a *one-sided bounded-error*, where either each $w \in P_{\text{yes}}$ is accepted with probability 1 or each $w \in P_{\text{no}}$ is rejected with probability 1.

If \mathcal{P} is a Las Vegas probabilistic finite automaton, we say that \mathcal{P} solves P with a *success probability* $p \in (0, 1]$, if (i) for each $w \in P_{\text{yes}}$, \mathcal{P} accepts w with probability at least p and never rejects (i.e., the probability of rejection is 0) and (ii) for each $w \in P_{\text{no}}$, \mathcal{P} rejects w with probability at least p and never accepts. (In both these cases, the remaining probability goes to “don’t know” answers.)

3 Basic Facts on Classical Automata

We continue with some basic facts regarding classical automata. We start with a basic observation: The cardinality of the class of the promise problems defined by 1DFAs is uncountable. We can show this fact even for the class defined by a fixed 2-state 1DFA on unary languages. Let $L \subseteq \mathcal{N}$ be any language defined on natural numbers. Then, the following unary promise problem $P(L)$ can be solvable by a 2-state 1DFA, say \mathcal{D} :

$$\begin{aligned} P_{\text{yes}}(L) &= \{a^{2n} \mid n \in L\} \\ P_{\text{no}}(L) &= \{a^{2n+1} \mid n \in L\} \end{aligned} ,$$

where \mathcal{D} has two states, the initial and only accepting state s_1 and s_2 , and, the automaton alternates between s_1 and s_2 for each input symbol. There is a one-to-one correspondence between $\{L \mid L \subseteq \mathcal{N}\}$ and $\{P(L) \mid L \subseteq \mathcal{N}\}$. Therefore, the cardinality of the class of the promise problems solvable by \mathcal{D} is uncountable. Moreover, if L is an uncomputable language, then both $P_{\text{yes}}(L)$ and $P_{\text{no}}(L)$ are uncomputable. The interested reader can easily find some other promise problems (solvable by very small 1DFAs) whose yes- and no-instances satisfy some special properties (i.e. non-context-free but context-sensitive, NP-complete, semi-decidable, etc.). One trivial construction for a binary language $L \in \{a, b\}^*$ is to define the promise problem as $P(L) = (\{aw \mid w \in L\}, \{bw \mid w \in L\})$. (A promise problem solvable by 1DFAs whose yes- and no-instances are non-regular languages previously given in [18].) Despite of the above facts, all trade-offs established between various models of automata stay valid also for promise problems.

First, we show that the classes of promise problems solvable by deterministic, nondeterministic, alternating, and Las Vegas probabilistic finite automata are identical.

Theorem 3.1 *If a promise problem $P = (P_{\text{yes}}, P_{\text{no}})$ can be solved by a 2AFA \mathcal{A} with n states, it can also be solved by a 1DFA \mathcal{D} with $t(n) \leq 2^{n \cdot 2^n}$ states.*

Proof: Let R denote the regular language recognized by \mathcal{A} . (That is, even though we are given a 2AFA \mathcal{A} for solving a promise problem P , such machine is still associated with some classical regular language $L(\mathcal{A}) = R$.) The complement of this language will be denoted by \bar{R} . Since the problem P is solvable by \mathcal{A} , we have that $P_{\text{yes}} \subseteq R$ and $P_{\text{no}} \subseteq \bar{R}$.

Now, by [27], each 2AFA \mathcal{A} with n states can be transformed into an equivalent 1DFA \mathcal{D} with $t(n) \leq 2^{n \cdot 2^n}$ states. This gives that $P_{\text{yes}} \subseteq R = L(\mathcal{D})$ and $P_{\text{no}} \subseteq \bar{R} = \overline{L(\mathcal{D})}$, that is, \mathcal{D} accepts each $w \in P_{\text{yes}}$ and rejects each $w \in P_{\text{no}}$. Therefore, \mathcal{D} can be used for solving P .

The proof of the above theorem can be easily updated for other classical models of automata — using the corresponding trade-off $t(n)$ for the conversion known from the literature — by which we obtain the following corollary.

Corollary 3.2 *Any trade-off $t(n)$ in the number of states for language recognition, between any two of deterministic, nondeterministic, or alternating automata models, is also a valid trade-off for promise problems.*

Next, we show that neither Las Vegas PFAs gain any computational power.

Theorem 3.3 *If a promise problem $P = (P_{\text{yes}}, P_{\text{no}})$ can be solved by a Las Vegas 2PFA \mathcal{P} with n states and any success probability $p > 0$, it can also be solved by a 1DFA \mathcal{D} with $t(n) \leq 2^{n^2+n}$ states.*

Proof: First, for each $w \in P_{\text{yes}}$, the machine \mathcal{P} has at least one accepting computation path and no rejecting path. Conversely, for each $w \in P_{\text{no}}$, the machine has at least one rejecting path and no accepting path. In both cases, the paths that are not successful return a “don’t know” answer.

Therefore, by removing the probabilities and converting neutral “don’t know” states into rejecting states, we obtain an n -state 2NFA \mathcal{N} recognizing a regular language R such that $P_{\text{yes}} \subseteq R$ and $P_{\text{no}} \cap R = \emptyset$. Therefore, $P_{\text{no}} \subseteq \bar{R}$.

Now, by [23], the 2NFA \mathcal{N} can be converted to an equivalent 1DFA \mathcal{D} with the number of states bounded by $t(n) \leq \sum_{i=0}^{n-1} \sum_{j=0}^{n-1} \binom{n-1}{i} \binom{n}{j} (2^i - 1)^j \leq 2^{n^2+n}$, which gives $P_{\text{yes}} \subseteq R = L(\mathcal{D})$ and $P_{\text{no}} \subseteq \bar{R} = \overline{L(\mathcal{D})}$. That is, \mathcal{D} accepts each $w \in P_{\text{yes}}$ and rejects each $w \in P_{\text{no}}$, and hence it can be used for solving P .

A similar idea can be used for Las Vegas 1_ϵ PFAs:

Corollary 3.4 *If a promise problem $P = (P_{\text{yes}}, P_{\text{no}})$ can be solved by a Las Vegas 1_ϵ PFA \mathcal{P} with n states and any success probability $p > 0$, it can also be solved by a 1DFA \mathcal{D} with $t(n) \leq 1 + 3^{(n-1)/3} \leq 2^{0.529 \cdot n}$ states.*

Proof: By removing the probabilities, we first obtain an n -state 1_ϵ SVFA \mathcal{N} recognizing a regular language R such that $P_{\text{yes}} \subseteq R$ and $P_{\text{no}} \subseteq \bar{R}$. Then, after elimination of ϵ -transitions by the method described in [26, Sect. 2.11] (still not increasing the number of states), we have a 1SVFA that can be converted to an equivalent 1DFA \mathcal{D} with $t(n) \leq 1 + 3^{(n-1)/3}$ states [22]. The resulting 1DFA is then used for solving P .

In Section 5, we show that the corresponding lower bound is also exponential.

4 A Classically Expensive Unary Promise Problem

Recently, Ambainis and Yakaryilmaz [2] presented a family of promise problems $\text{EvenOdd}(k)$, for $k \in \mathcal{N}_+$, where

$$\begin{aligned} \text{EvenOdd}_{\text{yes}}(k) &= \{a^{m \cdot 2^k} \mid m \text{ is even}\}, \\ \text{EvenOdd}_{\text{no}}(k) &= \{a^{m \cdot 2^k} \mid m \text{ is odd}\}, \end{aligned}$$

such that, for each k , the promise problem $\text{EvenOdd}(k)$ can be solved exactly by only a two-state realtime quantum finite automaton. On the other hand, for 1DFAs, 2^{k+1} states are necessary and also sufficient. Later, it was shown that 2^{k+1} states are also necessary in any bounded-error 1PFA [33] and that 2DFAs must use at least $2^{k+1} - 1$ states [3].⁵

Here we show that even two-way nondeterminism does not help us to save a single state for this promise problem; and hence 2^{k+1} states is the *exact* value for one-way/two-way deterministic/nondeterministic automata. After that, we show that alternation does help to reduce the number of states to $\Theta(k)$, already by the use of one-way ϵ -free machines. More precisely, $\text{EvenOdd}(k)$ can be solved by a 1AFA with $11k - 14$ states (for $k \geq 3$); the corresponding lower bound for 1AFAs is $k + 1$. For two-way machines, 2AFAs, the exact lower bound is an open problem, but we know that it must grow in k , which is an easy consequence of Theorems 3.1 and 4.1. Similarly, we do not know whether bounded-error 2PFAs can work with fewer than 2^{k+1} states.

⁵Recently [1], it was also shown that the width of any nondeterministic or bounded-error stable OBBD cannot be smaller than 2^{k+1} , for a functional version of $\text{EvenOdd}(k)$.

The method used for proving the lower bound in the following theorem is not new — inspired by [36, 6, 12] — but it needs some modifications on promise problems.

Theorem 4.1 *For each $k \in \mathcal{N}_+$, any 2NFA for solving the promise problem $\text{EvenOdd}(k)$ needs at least 2^{k+1} states.*

Proof: Suppose, for contradiction, that the promise problem can be solved by a 2NFA \mathcal{N} with $\|S\| < 2^{k+1}$ states.

Consider now the input $a^{2^{k+1}}$. Clearly, it must be accepted by \mathcal{N} , and hence there must exist at least one accepting computation path, so we can fix one such path. (Besides, being nondeterministic, \mathcal{N} can have also other paths, not all of them necessarily accepting, but we do not care for them.) Along this fixed path, take the sequence of states

$$q_0, p_1, q_1, p_2, q_2, p_3, q_3, \dots, p_r, q_r, p_{r+1},$$

where q_0 is the initial state, p_{r+1} is an accepting state, and all the other states $\{p_i \cup q_i \mid 1 \leq i \leq r\}$ are at the left/right endmarkers of the input 2^{k+1} , such that:

- If p_i is at the left (resp., right) endmarker, then q_i is at the opposite endmarker, and the path from p_i to q_i traverses the entire input from left to right (resp., right to left), not visiting any of endmarkers in the meantime.
- The path between q_i and p_{i+1} starts and ends at the same endmarker, possibly visiting this endmarker several times, but not visiting the opposite endmarker. Such a path is called a “U-turn”. This covers also the case of $q_i = p_{i+1}$ with zero number of executed steps in between.

Now, the path connecting p_i with q_i must visit all input tape positions in the middle of the input $a^{2^{k+1}}$, and hence \mathcal{N} must enter 2^{k+1} states, at least one state per each input tape position. However, since $\|S\| < 2^{k+1}$, some state r_i must be repeated. That is, between p_i and q_i , there must exist a loop, starting and ending in the same state r_i , traveling some ℓ_i positions to the right, not visiting any of the endmarkers. For each traversal, we fix one such loop (even though the path from p_i to q_i may contain many such loops). Note that the argument for traveling across the input from right to left is symmetrical. Since this loop is in the middle of the traversal, we have that $\ell_i < 2^{k+1}$. But then ℓ_i can be expressed in the form

$$\ell_i = \gamma_i \cdot 2^{\alpha_i},$$

where $\gamma_i > 0$ is odd (not excluding $\gamma_i = 1$) and $\alpha_i \in \{0, 1, \dots, k\}$. Note that if $\alpha_i \geq k+1$, we would have a contradiction with $\ell_i < 2^{k+1}$.

Now, consider the value

$$\ell = \text{lcm}\{\ell_1, \ell_2, \dots, \ell_r\},$$

the least common multiple of all fixed loops, for all traversals of the input $a^{2^{k+1}}$. Clearly, ℓ can also be expressed in the form

$$\ell = \gamma \cdot 2^\alpha,$$

where $\gamma > 0$ is odd and $\alpha \in \{0, 1, \dots, k\}$ (actually, $\gamma = \text{lcm}\{\gamma_1, \dots, \gamma_r\}$ and $\alpha = \max\{\alpha_1, \dots, \alpha_r\}$).

We shall now show that the machine \mathcal{N} must also accept the input

$$a^{2^{k+1} + \ell \cdot 2^{k-\alpha}}.$$

First, if q_i and p_{i+1} are connected by a U-turn path on the input $a^{2^{k+1}}$, they will be connected also on the input $a^{2^{k+1} + \ell \cdot 2^{k-\alpha}}$ since such path does not visit the opposite endmarker and $2^{k+1} + \ell \cdot 2^{k-\alpha} \geq 2^{k+1}$. Second, if p_i and q_i are connected by a left-to-right traversal along $a^{2^{k+1}}$, they

will stay connected also along the input $a^{2^{k+1}+\ell \cdot 2^{k-\alpha}}$. This only requires to repeat the loop of the length ℓ_i beginning and ending in the state r_i . Namely, we can make $\frac{\ell}{\ell_i} \cdot 2^{k-\alpha}$ more iterations. Note that $\ell = \text{lcm}\{\ell_1, \dots, \ell_r\}$ and hence ℓ is divisible by ℓ_i , for each $i = 1, \dots, r$, and that $\alpha \leq k$, which gives that $\frac{\ell}{\ell_i} \cdot 2^{k-\alpha}$ is an integer. Note also that these $\frac{\ell}{\ell_i} \cdot 2^{k-\alpha}$ additional iterations of the loop of the length ℓ_i travel exactly $(\frac{\ell}{\ell_i} \cdot 2^{k-\alpha}) \times \ell_i = \ell \cdot 2^{k-\alpha}$ additional positions to the right. Thus, if \mathcal{N} has an accepting path for $a^{2^{k+1}}$, it must also have an accepting path for $a^{2^{k+1}+\ell \cdot 2^{k-\alpha}}$ (just by a straightforward induction on r). Therefore, \mathcal{N} accepts $a^{2^{k+1}+\ell \cdot 2^{k-\alpha}}$. (Actually, \mathcal{N} can have many other paths for this longer input, but they cannot rule out the accepting decision of the path constructed above.)

However, the input $a^{2^{k+1}+\ell \cdot 2^{k-\alpha}}$ should be rejected, since $2^{k+1} + \ell \cdot 2^{k-\alpha} = 2^{k+1} + \gamma \cdot 2^\alpha \cdot 2^{k-\alpha} = (2+\gamma) \cdot 2^k$, where γ is odd. This is a contradiction. So, \mathcal{N} must have at least 2^{k+1} states.

The above argument cannot be extended to alternating automata: using “cooperation” of several computation paths running in parallel, the number of states can be reduced from 2^{k+1} to $O(k)$, even by the use of one-way ε -free machines. We first present a conceptually simpler realtime version.

Lemma 4.2 *For each $k \in \mathcal{N}_+$, the language $\text{EvenOdd}_{\text{yes}}(k)$ can be recognized — and hence the promise problem $\text{EvenOdd}(k)$ can be solved — by an rtAFA \mathcal{A} with $7k+2$ states.*

Proof: Recall that $\text{EvenOdd}_{\text{yes}}(k) = \{a^{m \cdot 2^k} \mid m \text{ is even}\}$, and hence the unary input string a^n is in this language if and only if n is an integer multiple of 2^{k+1} . In turn, this holds if and only if $b_{n,k} = b_{n,k-1} = \dots = b_{n,0} = 0$, where $b_{n,i}$ denotes the i^{th} bit in the binary representation of the number n . (The bit positions are enumerated from right to left, beginning with zero for the least significant bit.) That is, written in binary, n should be of the form $\gamma 0^{k+1}$, for some $\gamma \in \{0, 1\}^*$.

Based on this, the rtAFA \mathcal{A} counts the length of the input modulo 2^{k+1} and verifies, for the binary representation of this length, that the last $k+1$ bits are all equal to zero. The important detail is that \mathcal{A} counts *down*, starting from zero. That is, along the input, the value of the counter changes as follows:

$$\ell := 0, 2^{k+1}-1, 2^{k+1}-2, \dots, 2, 1, 0, 2^{k+1}-1, \dots$$

Thus, the counter ℓ is interpreted as the number of input symbols not processed yet (and taken modulo 2^{k+1}). The machine \mathcal{A} accepts if and only if this value is zero when it halts at the end of the input.

Implemented deterministically, this counting requires 2^{k+1} states. However, an alternating automaton can use several processes running in parallel. (See [34] for a similar idea.) Each parallel process will keep only *one bit* of the current value of ℓ . That is, for $i \in \{k, k-1, \dots, 0\}$, we shall use the states “ i_0 ” and “ i_1 ” to indicate that i^{th} bit in the current value of ℓ is cleared to zero or set to one, respectively. This will hold for each “proper” sequence of existential guesses along each input. The correctness of this guessing will be verified when \mathcal{A} halts at the end of the input. (Actually, the alternating tree of computation paths will be highly redundant; the same bit will be remembered by a huge number of activated identical copies of i_0 or i_1 , running in parallel.) Clearly, for each symbol a along the input, we need to decrement the counter ℓ by one and hence to update properly the corresponding bits. This is implemented by the use of some additional auxiliary states.

We are now ready to present the set of states for the rtAFA \mathcal{A} , together with their brief description:

i_x , for $i = k, \dots, 0$ and $x \in \{0, 1\}$: the i^{th} bit of ℓ is x ($b_{\ell,i} = x$),

i_{x,\forall_1} , for $i = k, \dots, 1$ and $x \in \{0, 1\}$: the i^{th} bit of ℓ is x and, moreover, all bits to the right are set to 1 ($b_{\ell,i} = x$ and $b_{\ell,j} = 1$ for each $j < i$),

i_{x,\exists_0} , for $i = k, \dots, 1$ and $x \in \{0, 1\}$: the i^{th} bit of ℓ is x and, moreover, at least one bit to the right is set to 0 ($b_{\ell,i} = x$ and $b_{\ell,j} = 0$ for some $j < i$),

i_{\exists_0} , for $i = k, \dots, 2$: the i^{th} bit of ℓ may be quite arbitrary, but at least one bit to the right is set to 0 ($b_{\ell,j} = 0$ for some $j < i$),

s_I : the initial state.

All states of type i_0, i_1, i_{\exists_0} are existential, whereas those of type $i_{0,\forall_1}, i_{1,\forall_1}, i_{0,\exists_0}, i_{1,\exists_0}, s_I$ are universal. The states of type i_0, s_I are accepting and none of the remaining states is accepting.

Clearly, the total number of states is $2(k+1) + 2k + 2k + (k-1) + 1 = 7k+2$.

Transitions in \mathcal{A} are defined as follows:

$i_x \xrightarrow{a} i_{x,\exists_0} \vee i_{1-x,\forall_1}$, for $i = k, \dots, 1$ and $x \in \{0, 1\}$: The i^{th} bit of ℓ is set to x if and only if either (i) the i^{th} bit of $\ell' = \ell - 1$ is set to the same value x , provided that at least one bit to the right in ℓ' is set to zero, that is, $b_{\ell',j} = 0$ for some $j < i$, or (ii) the i^{th} bit of ℓ' is flipped to the complementary value $1-x$, provided that all bits the right in ℓ' are set to one, that is, $b_{\ell',j} = 1$ for each $j < i$. Therefore, in the state i_x , branching existentially while reading the next input symbol a , \mathcal{A} guesses between switching to i_{x,\exists_0} or to i_{1-x,\forall_1} .

$0_x \xrightarrow{a} 0_{1-x}$, for $x \in \{0, 1\}$, a special case of i_x for $i = 0$: Transitions for the rightmost bit are simplified, since there is no bit position with $j < i$ and hence we do not need to use $0_{0,\forall_1}, 0_{1,\forall_1}, 0_{0,\exists_0}, 0_{1,\exists_0}$. Therefore, in 0_x , the machine reads the next input symbol a deterministically, by switching to 0_{1-x} .

$i_{x,\forall_1} \xrightarrow{\varepsilon} i_x \wedge (i-1)_1 \wedge \dots \wedge 0_1$, for $i = k, \dots, 1$ and $x \in \{0, 1\}$: In the auxiliary state i_{x,\forall_1} , the machine \mathcal{A} has to verify the bit setting at the corresponding positions in ℓ' , namely, that $b_{\ell',i} = x$ and that $b_{\ell',j} = 1$ for each $j < i$. Thus, branching universally with ε -transitions for each bit position $j \leq i$, it splits into parallel processes $i_x, (i-1)_1, \dots, 0_1$.

$i_{x,\exists_0} \xrightarrow{\varepsilon} i_x \wedge i_{\exists_0}$, for $i = k, \dots, 2$ and $x \in \{0, 1\}$: The role of the auxiliary state i_{x,\exists_0} is similar to that of i_{x,\forall_1} , verifying the bit setting in ℓ' . This time we verify that (i) $b_{\ell',i} = x$ and that (ii) $b_{\ell',j} = 0$ for at least one $j < i$. In order not to mix universal and existential decisions, we delay the condition (ii) until the next computation step, and hence \mathcal{A} branches universally in i_{x,\exists_0} , splitting into two parallel processes i_x and i_{\exists_0} , by the use of ε -transitions.

$1_{x,\exists_0} \xrightarrow{\varepsilon} 1_x \wedge 0_0$, for $x \in \{0, 1\}$, a special case of i_{x,\exists_0} for $i = 1$: Transitions are simplified, since we can work without 1_{\exists_0} . Therefore, in $1_{x,\exists_0}$, the machine splits universally into 1_x and 0_0 , using ε -transitions.

$i_{\exists_0} \xrightarrow{\varepsilon} (i-1)_0 \vee \dots \vee 0_0$, for $i = k, \dots, 2$: In the auxiliary state i_{\exists_0} , we verify that $b_{\ell',j} = 0$ for at least one $j < i$. Therefore, branching existentially with ε -transitions, \mathcal{A} guesses which bit is set to zero, switching to j_0 , for some $j < i$.

$s_I \xrightarrow{a} k_{1,\forall_1}$: At the very beginning, $\ell = n$. Therefore, for each accepted input a^n , the initial value of ℓ must be an integer multiple of 2^{k+1} , with the last $k+1$ bits all equal to zero. Thus, after reading the first input symbol a , we should get $\ell' = \ell - 1$ in which all these bits are set to one. For this reason, in the initial state, the machine reads the first input symbol a by switching deterministically to k_{1,\forall_1} .

It is easy to see from the above transitions that \mathcal{A} is a realtime machine: After reading an input symbol, a computation path can pass through at most two ε -transitions in a row, after which it gets to a state i_x , for some $i = k, \dots, 0$ and some $x \in \{0, 1\}$. In i_x , the machine waits for a next input symbol.

To show that \mathcal{A} indeed recognizes $\text{EvenOdd}_{\text{yes}}(k)$, we need to prove the following claim.

Claim. *For each $\ell \geq 0$, each $i = k, \dots, 0$, and each $x \in \{0, 1\}$, the rtAFA \mathcal{A} has an accepting alternating subtree of computation paths rooted in the state i_x at the input position ℓ (measured from the end of the input) if and only if the i^{th} bit of ℓ is equal to x .*

The claim is proved by induction on $\ell = 0, 1, 2, \dots$

First, let $\ell = 0$, that is, \mathcal{A} has already processed all input symbols. Clearly, all bits of ℓ are equal to zero, i.e., $b_{\ell,i} = 0$ for each i . Now, since there are no more input symbols to be processed and there are no ε -transitions starting in the state i_x , the machine \mathcal{A} must halt in i_x . Therefore, \mathcal{A} has an accepting alternating subtree rooted in i_x at the end of the input if and only if i_x is an accepting state. By definition of accepting states in \mathcal{A} (the states of type i_0 are accepting but those of type i_1 are not), this holds if and only if $x = 0$, which in turn holds if and only if $x = 0 = b_{\ell,i}$, the i^{th} bit of ℓ .

Second, by induction, assume that the claim holds for $\ell' = \ell - 1$. Now, by definition of the transitions in the state i_x for $i > 0$, taking into account subsequent ε -transitions passing through auxiliary states, it can be seen⁶ that \mathcal{A} has an accepting alternating subtree of computation paths rooted in the state i_x at the input position ℓ if and only if either (i) \mathcal{A} has an accepting alternating subtree rooted in the state i_x at the input position $\ell' = \ell - 1$ and, moreover, the same holds for at least one subtree rooted in a state j_0 at the position ℓ' , for some $j < i$, or (ii) \mathcal{A} has an accepting alternating subtree rooted in i_{1-x} at the input position ℓ' and, moreover, the same holds for all subtrees rooted in j_1 at the position ℓ' , for each $j < i$. However, by induction, using $\ell' = \ell - 1$, this statement holds if and only if either (i) $b_{\ell',i} = x$ and, moreover, $b_{\ell',j} = 0$, for some $j < i$, or (ii) $b_{\ell',i} = 1 - x$ and, moreover, $b_{\ell',j} = 1$, for each $j < i$. Using $\ell = \ell' + 1$, it is easy to see that this holds if and only if $b_{\ell,i} = x$.

This completes the proof of the claim.

Now, recall that the machine \mathcal{A} reads the first input symbol a by switching deterministically from the initial state to the state k_{1,\vee_1} . After that, branching universally with ε -transitions for each bit position $j \leq k$, the computation splits into parallel processes $k_1, (k-1)_1, \dots, 0_1$. Thus, using the above claim, \mathcal{A} accepts a^n with $n > 0$ if and only if, in $\ell = n - 1$, the last $k + 1$ bits are all equal to 1. Clearly, this holds if and only if the last $k + 1$ bits in n are all equal to 0, that is, if and only if n is an integer multiple of 2^{k+1} . This gives that the language recognized by \mathcal{A} agrees with $\text{EvenOdd}_{\text{yes}}(k)$ on all inputs of length $n > 0$.

For the special case of $n = 0$, we have that $a^0 = a^{0 \cdot 2^k} \in \text{EvenOdd}_{\text{yes}}(k)$, since $m = 0$ is even. However, $a^0 = \varepsilon$ is accepted by \mathcal{A} due to the fact that the initial state s_1 is also an accepting state, which completes the argument.

By increasing slightly the number of states, we can replace the realtime alternating machine from the above lemma by a one-way AFA not using any ε -transitions:

Theorem 4.3 *For each $k \geq 3$, the language $\text{EvenOdd}_{\text{yes}}(k)$ can be recognized — and hence the promise problem $\text{EvenOdd}(k)$ can be solved — by a 1AFA \mathcal{A}'' with $11k - 14$ states.*

Proof: By inspection of the rtAFA \mathcal{A} constructed in Lemma 4.2, it can be seen that \mathcal{A} can read an input symbol only by transitions starting in the states of type i_x, s_1 . After reading, the

⁶With obvious modifications, the argument presented here can be easily extended also for the case of $i = 0$ (the rightmost bit), which we leave to the reader.

computation path can pass through at most two subsequent ε -transitions until it gets to a state of type i_x , ready to read from the input again.

First, as an intermediate product, we replace \mathcal{A} by an equivalent machine \mathcal{A}' in which reading an input symbol is always followed by *exactly three* subsequent ε -transitions, along each computation path. This only requires to delay artificially the moment when the automaton gets to a state of type i_x by the use of some additional states, namely, the following additional states:

i'_x, i''_x , for $i = k, \dots, 0$ and $x \in \{0, 1\}$: delay switching to the state i_x by one or two steps, respectively — implemented by transitions $i''_x \xrightarrow{\varepsilon} i'_x \xrightarrow{\varepsilon} i_x$,

$0'''_x$, for $x \in \{0, 1\}$: delay switching to the state 0_x by three steps — implemented by $0'''_x \xrightarrow{\varepsilon} 0''_x$.

Taking also into account the states used by the original machine \mathcal{A} , the total number of states increases to $(7k+2) + 4(k+1) + 2 = 11k+8$.

Transitions for the “original” states (cf. construction in Lemma 4.2) are modified as follows:

$$i_x \xrightarrow{a} i_{x,\exists_0} \vee i_{1-x,\forall_1}, \text{ for } i = k, \dots, 1 \text{ and } x \in \{0, 1\} \quad (\text{not modified}),$$

$$0_x \xrightarrow{a} 0'''_{1-x}, \text{ for } x \in \{0, 1\},$$

$$i_{x,\forall_1} \xrightarrow{\varepsilon} i''_x \wedge (i-1)''_1 \wedge \dots \wedge 0''_1, \text{ for } i = k, \dots, 1 \text{ and } x \in \{0, 1\},$$

$$i_{x,\exists_0} \xrightarrow{\varepsilon} i''_x \wedge i_{\exists_0}, \text{ for } i = k, \dots, 2 \text{ and } x \in \{0, 1\},$$

$$1_{x,\exists_0} \xrightarrow{\varepsilon} 1''_x \wedge 0''_0, \text{ for } x \in \{0, 1\},$$

$$i_{\exists_0} \xrightarrow{\varepsilon} (i-1)'_0 \vee \dots \vee 0'_0, \text{ for } i = k, \dots, 2,$$

$$s_1 \xrightarrow{a} k_{1,\forall_1} \quad (\text{not modified}).$$

Since \mathcal{A}' and \mathcal{A} share the same set of accepting states, it should be obvious that they recognize the same language.

Second, replace all ε -transitions in \mathcal{A}' by transitions reading the input symbol a . That is, each transition of type $s_1 \xrightarrow{\varepsilon} s_2$ is replaced by $s_1 \xrightarrow{a} s_2$; the transitions already reading input symbols do not change; nor does the set of accepting states. This way we obtain a one-way AFA \mathcal{A}'' without any ε -transitions such that, for each $n \geq 0$, it accepts the input a^{4n} if and only if \mathcal{A}' accepts a^n , i.e., if and only if $a^n \in \text{EvenOdd}_{\text{yes}}(k)$. The new machine \mathcal{A}'' does not accept any input the length of which is not an integer multiple of 4.

Therefore, with $11k+8$ states, \mathcal{A}'' recognizes the language

$$\{(a^4)^{m \cdot 2^k} \mid m \text{ is even}\} = \{a^{m \cdot 2^{k+2}} \mid m \text{ is even}\} = \text{EvenOdd}_{\text{yes}}(k+2).$$

Using the substitution $k' = k+2$, we thus get that $\text{EvenOdd}_{\text{yes}}(k')$ can be recognized by a one-way ε -free AFA \mathcal{A}'' with $11(k'-2) + 8 = 11k' - 14$ states, for each $k' \geq 3$.

To show that the gap exhibited by Theorems 4.1 and 4.3 is tight, we need an upper bound for converting alternation into determinism on unary promise problems:

Lemma 4.4 *If a promise problem $P = (P_{\text{yes}}, P_{\text{no}})$ built over a unary input alphabet can be solved by a 1AFA with n states, it can also be solved by a 1DFA with $t(n) \leq 2^n$ states.*

Proof: As pointed out in [28, Sect. 5], each 1AFA \mathcal{A} with n states recognizing a unary language U can be replaced by an equivalent 1DFA \mathcal{D} with 2^n states. This is a simple consequence of the fact that \mathcal{A} can be replaced by a 1DFA \mathcal{D} with 2^n states recognizing U^R , the reversal of the original language [5]. (Actually, the construction in [5] works for Boolean automata — see Footnote 3 — which covers, as a special case, the classical alternating automata as well.) However, for unary languages, we have $U^R = U$.

Then, by the same reasoning as used in the proof of Theorem 3.1, this trade-off easily extends to unary promise problems.

By the above lemma, the lower bound for solving $\text{EvenOdd}(k)$ by 1AFAS is $k+1$ states, since a 1AFA with less than $k+1$ states would give us a 1DFA with less than 2^{k+1} states for this promise problem, which contradicts Theorem 4.1.

By combining Theorems 4.1 and 4.3 with Lemma 4.4, we thus get:

Corollary 4.5 *The tight gap of succinctness between two-way NFAs and one-way ε -free AFAs is asymptotically exponential on promise problems, even for a unary input alphabet.*

5 Las Vegas Probabilistic Finite Automata

In the case of language recognition, the gap of succinctness between Las Vegas realtime PFAs and one-way DFAs can be at most quadratic and this gap is tight [8, 21]. Quite surprisingly, despite of this fact and despite of the fact that the existing gaps for language recognition are not be violated on promise problems by classical models of automata (cf. Corollary 3.2), we shall show here that this does not hold for the case of one-way Las Vegas probabilistic versus one-way deterministic automata; by switching from language recognition to promise problems, the quadratic gap increases to a tight exponential gap.

For this purpose, we introduce a family of promise problems solvable by Las Vegas 1PFAs with a linear number of states and success probability arbitrarily close to 1, but for which any 1DFAs require an exponential number of states.⁷

Consider a family of promise problems $\text{TRIOS}(n, r)$, for $n, r \in \mathcal{N}_+$, such that

$$\begin{aligned} \text{TRIOS}_{\text{yes}}(n, r) &= \{ \#x_1x_1y_1\#x_2x_2y_2\cdots\#x_rx_ry_r \mid x_1, y_1, \dots, x_r, y_r \in \{0, 1\}^n, \\ &\quad \forall i \in \{1, \dots, r\} \exists j \in \{1, \dots, n\} : x_{i,j} < y_{i,j} \}, \\ \text{TRIOS}_{\text{no}}(n, r) &= \{ \#x_1y_1x_1\#x_2y_2x_2\cdots\#x_rx_ry_r \mid x_1, y_1, \dots, x_r, y_r \in \{0, 1\}^n, \\ &\quad \forall i \in \{1, \dots, r\} \exists j \in \{1, \dots, n\} : x_{i,j} > y_{i,j} \}, \end{aligned}$$

where $x_{i,j}, y_{i,j} \in \{0, 1\}$ denote the j^{th} bit in the respective strings $x_i, y_i \in \{0, 1\}^n$.

Theorem 5.1 *For each $n, r \in \mathcal{N}_+$, the promise problem $\text{TRIOS}(n, r)$ can be solved by a Las Vegas 1PFA \mathcal{P} with $4n+3$ states and success probability $1 - (\frac{n-1}{n})^r$.*

Proof: The automaton \mathcal{P} uses the following state set:

$$Q = \{q_I, q_A, q_R\} \cup \{s_i \mid i \in \{1, \dots, n\}\} \cup \{t_{i,0} \mid i \in \{1, \dots, 2n\}\} \cup \{t_{i,1} \mid i \in \{1, \dots, n\}\},$$

where q_I is the initial and also the “don’t know” state, while q_A, q_R are the accepting and rejecting states, respectively.

Recall that we do not have to worry about the outcome of the computation if the input is not of the form $(\#\{0, 1\}^{3n})^r$. Now, let $\#x_iu_iv_i$, with $x_i, u_i, v_i \in \{0, 1\}^n$, denote the i^{th} input segment. Typically, for $i > 1$, the processing of the i^{th} segment is activated already somewhere in the middle of the previous segment, by switching to the initial state q_I . For $i = 1$, the machine begins with the first symbol $\#$. For these reasons, in the state q_I , the machine ignores all symbols along the input until it reads the separator $\#$, when it enters the state s_1 . This is implemented by transitions $q_I \xrightarrow{b} q_I$, for each $b \in \{0, 1\}$, together with $q_I \xrightarrow{\#} s_1$. (Unless otherwise stated explicitly, all transitions throughout this proof are “deterministic”, i.e., with probability 1.)

⁷Remark that, in [15], we argued the same exponential separation by using some other families of promise problems, but it appeared that some of these families can also be solved by 1DFAs with a linear number of states, and hence they cannot be used as witness families for proving such separation. More precisely, a single sentence statement given just before Corollary 4 in [15] is not correct.

Next, starting in the state s_1 , the automaton \mathcal{P} moves along the string x_i and chooses, with equal probability $\frac{1}{n}$, one of the symbols $x_{i,1}, \dots, x_{i,n} \in \{0, 1\}$ for subsequent verification. This is implemented as follows. If \mathcal{P} is in the state s_j , for $j \in \{1, \dots, n\}$, and the next input symbol is $b \in \{0, 1\}$, then \mathcal{P} reads this symbol by switching to the state $t_{1,b}$ with probability $p_j = \frac{1}{p'_1 \cdots p'_{j-1}} \cdot \frac{1}{n}$ and to the state s_{j+1} with probability $p'_j = 1 - p_j$. (As usual in mathematical notation, we take $p'_1 \cdots p'_{j-1} = 1$ for $j = 1$.) The reader may easily verify that the sequence of probabilities p_1, \dots, p_n ends with $p_n = 1$, which gives $p'_n = 0$, and hence there is no need to introduce a state s_{n+1} . It is also easy to see that \mathcal{P} gets from the state s_1 at the beginning of $x_i = x_{i,1} \cdots x_{i,n}$ to the state $t_{1,x_{i,j}}$ positioned just behind the symbol $x_{i,j}$ with probability $\frac{1}{n}$.

The next phase depends on whether \mathcal{P} is in the state $t_{1,0}$ or in $t_{1,1}$, that is, on whether the chosen bit $x_{i,j}$ is equal to 0 or to 1.

For the case of $x_{i,j} = 0$, the routine starting in the state $t_{1,0}$ is responsible to verify that $v_{i,j} = 1$, which corresponds to the case of $\#x_i u_i v_i = \#x_i x_i y_i$, with $x_{i,j} < y_{i,j}$. Since \mathcal{P} gets to $t_{1,0}$ just after reading the symbol $x_{i,j}$, this only requires to examine the bit that is placed exactly $2n$ positions to the right. This is implemented by transitions $t_{j,0} \xrightarrow{b} t_{j+1,0}$, for each $j \in \{1, \dots, 2n-1\}$ and each $b \in \{0, 1\}$. Finally, if $v_{i,j} = 1$, then \mathcal{P} accepts, using $t_{2n,0} \xrightarrow{1} q_A$. Otherwise, \mathcal{P} switches back to the initial state, by $t_{2n,0} \xrightarrow{0} q_I$. In the accepting state q_A , the machine just consumes the remaining part of the input, using $q_A \xrightarrow{a} q_A$, for each $a \in \{0, 1, \#\}$. Conversely, in the initial state q_I , the machine proceeds to examine the next input segment $\#x_{i+1} u_{i+1} v_{i+1}$, by transitions already described above. However, if \mathcal{P} switches to q_I while processing the last input segment (for $i = r$), it halts in q_I at the end of the input. This is interpreted as a “don’t know” answer.

The case of $x_{i,j} = 1$ is very similar. This time the routine starting in the state $t_{1,1}$ verifies that $u_{i,j} = 0$, which corresponds to the case of $\#x_i u_i v_i = \#x_i y_i x_i$, with $x_{i,j} > y_{i,j}$. This only requires to examine the bit placed exactly n positions to the right of the symbol $x_{i,j}$. This is implemented by transitions $t_{j,1} \xrightarrow{b} t_{j+1,1}$, for each $j \in \{1, \dots, n-1\}$ and each $b \in \{0, 1\}$. After that, if $u_{i,j} = 0$, then \mathcal{P} rejects, using $t_{n,1} \xrightarrow{0} q_R$. Otherwise, \mathcal{P} switches back to the initial state, by $t_{n,1} \xrightarrow{1} q_I$. In the rejecting state q_R , the machine consumes the rest of the input, using $q_R \xrightarrow{a} q_R$, for each $a \in \{0, 1, \#\}$. Conversely, in the initial state q_I , the machine proceeds in the same way as in the case of $x_{i,j} = 0$; either it proceeds to the next input segment, or else it halts in q_I at the end of the input, giving a “don’t know” answer.

Summing up, if the i^{th} segment is of the form $\#x_i u_i v_i = \#x_i x_i y_i$, with $x_{i,j} < y_{i,j}$ for some $j \in \{1, \dots, n\}$, then \mathcal{P} , starting to process this segment in the state q_I , accepts the entire input with probability at least $\frac{1}{n}$ (the exact value increases in the number of bit pairs satisfying $x_{i,j} < y_{i,j}$) and proceeds to the next segment in the state q_I with probability at most $1 - \frac{1}{n}$. (The machine never rejects on such segment, since $u_i = x_i$ and there are no bit pairs satisfying $x_{i,j} > u_{i,j}$.)

Now, recall that each $w \in \text{TRIOS}_{\text{yes}}(n, r)$ consists of r segments having this form, and hence \mathcal{P} gives the “don’t know” answer with probability at most $(1 - \frac{1}{n})^r$ for w . Therefore, w is accepted with probability at least $1 - (1 - \frac{1}{n})^r$ and never rejected.

Conversely, if i^{th} segment is of the form $\#x_i u_i v_i = \#x_i y_i x_i$, with $x_{i,j} > y_{i,j}$ for some $j \in \{1, \dots, n\}$, then, in the course of processing this segment, \mathcal{P} rejects the entire input with probability at least $\frac{1}{n}$ and proceeds to the next segment with probability at most $1 - \frac{1}{n}$. (Here the machine never accepts, since $v_i = x_i$ and there are no bit pairs with $x_{i,j} < v_{i,j}$.)

But then each $w \in \text{TRIOS}_{\text{no}}(n, r)$, consisting of r segments of this form, is rejected with probability at least $1 - (1 - \frac{1}{n})^r$ and never accepted, by the same reasoning as for $\text{TRIOS}_{\text{yes}}(n, r)$.

It is obvious that, for arbitrarily small fixed constant $\varepsilon > 0$, by taking $r = \lceil \log \varepsilon / \log(\frac{n-1}{n}) \rceil$, we shall obtain a family of promise problems, namely, $\text{TRIOS}'_\varepsilon(n)$, for $n \in \mathcal{N}_+$, where $\text{TRIOS}'_\varepsilon(n) = \text{TRIOS}(n, \lceil \log \varepsilon / \log(\frac{n-1}{n}) \rceil)$, such that the n^{th} member of this family can be solved by a Las Vegas 1PFA with $4n+3$ states and success probability at least $1 - \varepsilon$.

On the other hand, for each fixed $r \in \mathcal{N}_+$, 1DFAs need an exponential number of states for solving $\text{TRIOS}(n, r)$.

Theorem 5.2 *For each $n, r \in \mathcal{N}_+$, any 1DFA needs at least 2^n states for solving the promise problem $\text{TRIOS}(n, r)$.*

Proof: For contradiction, suppose that this promise problem can be solved by a 1DFA \mathcal{D} using $\|S\| < 2^n$ states. Let s_1 be its initial state, and let $\delta(q, u) \in S$ denote the state reached by \mathcal{D} from the state $q \in S$ by reading the string $u \in \Sigma^*$. Clearly, $\delta(q, uv) = \delta(\delta(q, u), v)$, for each state q and each two strings u, v .

Now, for each state $q \in S$, consider the enumeration of states $\delta(q, \#u)$, running over all strings $u \in \{0, 1\}^n$. Since the number of states in S is smaller than the number of strings in $\{0, 1\}^n$, we have, for each $q \in S$, that there must exist at least two different strings $u, v \in \{0, 1\}^n$ with $\delta(q, \#u) = \delta(q, \#v)$. Using some lexicographic ordering for $\{0, 1\}^n$, we can pick u_q and v_q , the first pair of strings satisfying $u_q \neq v_q$ and $\delta(q, \#u_q) = \delta(q, \#v_q)$. Moreover, without loss of generality, we can assume that, for some bit position $j \in \{1, \dots, n\}$, we have $u_{q,j} < v_{q,j}$. (Here $u_{q,j}, v_{q,j}$ denote the j^{th} bit in the respective strings u_q, v_q .) Otherwise, by renaming, the roles of u_q and v_q can be swapped.

Consider now the input

$$\begin{aligned} w_A &= \#u_{q_1}u_{q_1}v_{q_1}\#u_{q_2}u_{q_2}v_{q_2}\cdots\#u_{q_r}u_{q_r}v_{q_r}, \text{ where} \\ q_1 &= s_1, \\ q_{i+1} &= \delta(q_i, \#u_{q_i}u_{q_i}v_{q_i}), \text{ for } i = 1, \dots, r. \end{aligned}$$

It should be easily seen that, along the input w_A , \mathcal{D} passes the segment boundaries by the sequence of states $q_1, q_2, \dots, q_r, q_{r+1}$. More precisely, for each $i = 1, \dots, r$, the machine is in the state q_i at the beginning of the i^{th} segment $\#u_{q_i}u_{q_i}v_{q_i}$, halting in the state q_{r+1} at the end of the last segment. It is also obvious that $w_A \in \text{TRIOS}_{\text{yes}}(n, r)$, since, for each segment $\#u_{q_i}u_{q_i}v_{q_i}$, we have a bit position $j \in \{1, \dots, n\}$ satisfying $u_{q_i,j} < v_{q_i,j}$. Therefore, q_{r+1} is an accepting state.

Next, consider the input

$$w_R = \#v_{q_1}u_{q_1}v_{q_1}\#v_{q_2}u_{q_2}v_{q_2}\cdots\#v_{q_r}u_{q_r}v_{q_r}.$$

Clearly, also here the computation starts in the state $q_1 = s_1$. Now, using the fact that $\delta(q, \#u_q) = \delta(q, \#v_q)$ for each state $q \in S$, we have that q_{i+1} can also be expressed as

$$\begin{aligned} q_{i+1} &= \delta(q_i, \#u_{q_i}u_{q_i}v_{q_i}) = \delta(\delta(q_i, \#u_{q_i}), u_{q_i}v_{q_i}) \\ &= \delta(\delta(q_i, \#v_{q_i}), u_{q_i}v_{q_i}) = \delta(q_i, \#v_{q_i}u_{q_i}v_{q_i}), \end{aligned}$$

for each $i = 1, \dots, r$. Therefore, along the input w_R , \mathcal{D} passes the segment boundaries by the same sequence of states $q_1, q_2, \dots, q_r, q_{r+1}$, halting in q_{r+1} . But then, since q_{r+1} is an accepting state, \mathcal{D} accepts w_R . However, $w_R \in \text{TRIOS}_{\text{no}}(n, r)$, since, for each segment $\#v_{q_i}u_{q_i}v_{q_i}$, we have a bit position $j \in \{1, \dots, n\}$ satisfying $v_{q_i,j} > u_{q_i,j}$.

Thus, each 1DFA \mathcal{D} with less than 2^n states accepts an input that should be rejected, a contradiction.

As a consequence of the above theorem, we get the asymptotically optimal exponential size $\Theta(2^n)$ for 1DFAs solving $\text{TRIOS}(n, r)$; the construction of a 1DFA with $O(2^n)$ states is quite straightforward.

On the other hand, Theorem 5.1 guarantees $4n + 3$ states for Las Vegas 1PFAs solving $\text{TRIOS}(n, r)$, and hence we have an exponential lower bound for the gap between these two models of automata. Finally, by Corollary 3.4, we have an exponential upper bound, which gives:

Corollary 5.3 *The tight gap of succinctness between one-way DFAs and Las Vegas one-way ε -free PFAs is asymptotically exponential on promise problems.*

Due to Theorem 5.1, we know that $\text{TRIOS}(n, n^2)$ can be solved by a Las Vegas 1PFA with $4n+3$ states and success probability $\sigma(n) \geq 1 - (1 - \frac{1}{n})^{n^2}$. Thus, using the fact [20, Lm. A.3.60+] that $(1 - \frac{1}{x})^x < \frac{1}{e}$ for each real $x > 1$, the probability of a “don’t know” outcome is at most

$$1 - \sigma(n) \leq \left(1 - \frac{1}{n}\right)^{n \cdot n} < \left(\frac{1}{e}\right)^n = \frac{1}{e^n}.$$

This Las Vegas 1PFA can be converted to an equivalent restarting one-way PFA, executing the procedure in an infinite loop, by restarting the entire computation when the original machine halts in the “don’t know” state at the end of the input. (This can be easily achieved by using a single additional state.) Therefore, we can obtain a restarting 1PFA with $4n+4$ states, solving $\text{TRIOS}(n, n^2)$ exactly. The expected number of the input tape sweeps can be bounded by

$$\varrho(n) = \sum_{i=1}^{\infty} i \cdot \sigma(n) \cdot (1 - \sigma(n))^{i-1} \leq \sum_{i=1}^{\infty} i \cdot 1 \cdot \left(\frac{1}{e^n}\right)^{i-1} = \frac{1}{\left(1 - \frac{1}{e^n}\right)^2} = \left(1 + \frac{1}{e^n - 1}\right)^2 \leq 1 + o(1).$$

Thus, for the input of length $|w|$, the expected running time is $(1 + o(1)) \cdot |w|$.

Finally, the family of promise problems $\text{TRIOS}(n, r)$, for $n, r \in \mathcal{N}_+$, does not give corresponding separations for two-way machines, since already 2DFAs can solve $\text{TRIOS}(n, r)$ with $O(n)$ states, for each r . This is done as follows. Let $\#x_i u_i v_i$, with $x_i, u_i, v_i \in \{0, 1\}^n$, denote the i^{th} segment along the input. Processing of this segment begins with positioning the input head at the n^{th} bit of u_i , by moving $2n$ positions to the right from the symbol $\#$. Then, for $j = n, \dots, 1$ (in that order), the machine verifies that $u_{i,j} = x_{i,j}$. (To move from the j^{th} bit of u_i to the j^{th} bit of x_i , just travel exactly n positions to the left; to move from the j^{th} bit of x_i to the $(j-1)^{\text{st}}$ bit of u_i , just travel exactly $n-1$ positions to the right. The machine detects that it has tested all n bits from the fact that, after traveling exactly n positions to the left for the j^{th} bit of x_i , it finds the symbol $\#$ at this position, instead of zero or one.) In a similar way, the machine can also search for a bit position j satisfying $u_{i,j} < v_{i,j}$, this time iterating in ascending order, for $j = 1, \dots, n$. If the i^{th} input segment passes the test successfully, the machine proceeds to examine the next segment $\#x_{i+1}u_{i+1}v_{i+1}$ or, for $i = r$, it accepts.

6 Two-Sided Bounded-Error Probabilistic Finite Automata

In Section 3, we have shown the limitations of Las Vegas PFAs. One-sided error⁸ PFAs have similar limitations since they can be simulated by nondeterministic or universal finite automata by removing the probabilities: If a promise problem $P = (P_{\text{yes}}, P_{\text{no}})$ can be solved by a PFA rejecting each $w \in P_{\text{no}}$ with probability 1 (only members of P_{yes} are accepted), then we obtain an NFA from the one-sided error PFA. Conversely, if each $w \in P_{\text{yes}}$ is accepted with probability 1 (only members of P_{no} are rejected), then we obtain a finite automaton making only universal decisions (or an NFA for P_{no}) from the one-sided error PFA. In this section, we show that we have a different picture for PFAs with two-sided bounded-error.

First, we show that bounded-error 1PFAs can be very succinct compared to Las Vegas 2PFAs, 2AFAs, or any other simpler machines. In fact, we present a parallel result to that of Ambainis and Yakaryılmaz [2] but with bounded error instead of exact acceptance.

To this aim, consider \mathcal{U}_p , a 1PFA with the unary input alphabet $\Sigma = \{a\}$ and two states. Namely, let s_I be the initial and also the only accepting state, and s_R be the rejecting state. In the state s_I , the machine reads the symbol a staying in s_I with probability p and switching

⁸We also cover the case of *unbounded-error*, i.e., the machines with no bound on the error and so the error can be arbitrarily close to 1 for some inputs.

to s_R with the remaining probability $p' = 1-p$. In s_R , the machine just consumes the remaining part of the input, executing a single-state loop.

Now, let $\text{UP}(p)$ be a promise problem such that

$$\begin{aligned}\text{UP}_{\text{yes}}(p) &= \{a^j \mid f_{\mathcal{U}_p}(a^j) \geq \tfrac{3}{4}\}, \\ \text{UP}_{\text{no}}(p) &= \{a^j \mid f_{\mathcal{U}_p}(a^j) \leq \tfrac{1}{4}\}.\end{aligned}$$

Clearly, for each $p \in (0, 1)$, the promise problem $\text{UP}(p)$ can be solved by a bounded-error one-way ε -free PFA with only two states, namely, by \mathcal{U}_p .

On the other hand, it is easy to show that the number of states required by any 2AFA for solving $\text{UP}(p)$ (which covers all simpler models of automata as well) increases when p approaches 1: Since it is straightforward that the string a^j is accepted by \mathcal{U}_p with probability p^j for each $j \geq 0$, we get $a^j \in \text{UP}_{\text{yes}}(p)$ only for finitely many values satisfying $j \leq \lfloor \log_p \frac{3}{4} \rfloor$, whereas $\text{UP}_{\text{no}}(p)$ contains infinitely many strings satisfying $j \geq \lceil \log_p \frac{1}{4} \rceil$. Let us denote these two critical lengths as

$$A_p = \lfloor \log_p \tfrac{3}{4} \rfloor \quad \text{and} \quad R_p = \lceil \log_p \tfrac{1}{4} \rceil.$$

Note also that $A_p \rightarrow +\infty$ as $p \rightarrow 1$.

Theorem 6.1 *For one-way/two-way deterministic/nondeterministic finite automata, and each $p \in (0, 1)$, the number of states that is sufficient and necessary for solving the promise problem $\text{UP}(p)$ is exactly $A_p + 1$.*

Proof: First, we can easily design a 1DFA \mathcal{D} (hence, any more powerful model as well) solving \mathcal{U}_p with $A_p + 1$ states. After the initial chain of $A_p + 1$ accepting states, responsible to accept a^j for $j \in \{0, \dots, A_p\}$, the automaton enters a rejecting state s_R , in which it executes a single-step loop, consuming the rest of the input. By allowing undefined transitions in \mathcal{D} , the state s_R can be eliminated.

Second, before providing the matching lower bound, we need to present the so-called “ $n \rightarrow n+n!$ ” method [36, 6, 12] (cf. also the proof of Theorem 4.1). To make the paper self-contained, we include also a proof, for a simplified version of this method that is sufficient for our purposes.

Claim. *Let \mathcal{N} be a 2NFA (or any less powerful model) with n states and let a^m be a unary string, with $m \geq n$. Then, if \mathcal{N} has a computation path traversing across a^m from left to right, starting in some state q_1 and ending in some q_2 , it has also a computation path traversing across $a^{m+h \cdot m!}$, for each $h \geq 1$, starting and ending in the same states q_1 and q_2 .*

The proof is based on the fact that \mathcal{N} uses only n states, but the sequence of states along the left-to-right traversal across a^m is of length $m+1 > n$ (including q_1 at the beginning and q_2 at the end). Thus, in the course of this traversal, some state must be repeated. Hence, \mathcal{N} executes a loop, by which it travels some $\ell \leq m$ positions to the right. Note that $m!/\ell$ is a positive integer. Hence, for each $h \geq 1$, by using $h \cdot m!/\ell$ additional iterations of this loop traveling ℓ positions, the machine traverses the string $a^{m+h \cdot m!}$, starting and ending in the same states q_1 and q_2 . This completes the proof of the claim.

Now, let \mathcal{N} be a 2NFA (or any less powerful model) solving $\text{UP}(p)$ with at most A_p states. Consider the inputs a^m and $a^{m+h \cdot m!}$, for $m = A_p$ and arbitrary $h \geq 1$. Clearly, $a^m = a^{A_p} \in \text{UP}_{\text{yes}}(p)$, and hence \mathcal{N} has at least one accepting computation path for this input.

However, by the above claim, whenever this computation path traverses from the left endmarker of a^m to the opposite endmarker, starting in some state q_1 and ending in some q_2 , an updated version of this path can traverse across $a^{m+h \cdot m!}$ starting and ending in the same states q_1 and q_2 . By symmetry, the same holds for right-to-left traversals. Moreover, a U-turn starting and ending at the same endmarker of a^m can obviously do the same on the input $a^{m+h \cdot m!}$. Thus, for each $h \geq 1$, \mathcal{N} has a computation path accepting $a^{m+h \cdot m!}$. But then \mathcal{N} accepts infinitely many inputs, which contradicts the fact that all inputs longer than R_p belong to $\text{UP}_{\text{no}}(p)$ and should be rejected. Thus, \mathcal{N} must use at least $A_p + 1$ states.

A lower bound for 2AFAs solving the promise problem $\text{UP}(p)$ can be obtained by combining Corollary 3.2, Theorem 6.1, and the result of Geffert and Okhotin [14], stating that an arbitrary n -state 2AFA can be converted into an equivalent 1NFA with $2^{O(n \log n)}$ states. From this one can derive that the number of states in any 2AFA solving $\text{UP}(p)$ is at least $\Omega(\frac{\log A_p}{\log \log A_p})$.

Corollary 6.2 *There exists $\{\text{UP}(p) \mid p \in (0, 1)\}$, a family of unary promise problems such that bounded-error one-way ε -free PFAs with only two states are sufficient to solve all family, but the number of states required by 2AFAs (or by any simpler machines) to solve this family cannot be bounded by any constant.*

The error bound $\frac{1}{4}$ in the definition of $\text{UP}(p)$ can be replaced with an arbitrarily small error; the same results will follow by the use of the same reasoning.

Next, we present a separation result between bounded-error 1PFAs and 1DFAs (and so 2AFAs, or any other model capable of recognizing only regular languages). For this purpose, we use an idea given by Jibrán and Yakaryılmaz [33]. It is known that 2PFAs can recognize some nonregular languages, e.g., $\text{EQ} = \{a^n b^n \mid n \in \mathcal{N}_+\}$ [11], with bounded error but this requires exponential time [10]. It was also shown that EQ can be recognized by a restarting 1PFA for any error bound [38]. If the given string $a^m b^n$, where $m, n \in \mathcal{N}_+$, can be examined exponentially many times, then the algorithm given in [38] for EQ can distinguish between the cases of $m = n$ and $m \neq n$ with high probability. Based on this, we introduce a new family of promise problems $\text{ExpEQ}(c)$, for integer $c \geq 3$:

$$\begin{aligned} \text{ExpEQ}_{\text{yes}}(c) &= \{(a^m b^n)^{3 \cdot (2c^2)^{m+n} \cdot \lceil \ln c \rceil} \mid m, n \in \mathcal{N}_+, m = n\}, \\ \text{ExpEQ}_{\text{no}}(c) &= \{(a^m b^n)^{3 \cdot (2c^2)^{m+n} \cdot \lceil \ln c \rceil} \mid m, n \in \mathcal{N}_+, m \neq n\}. \end{aligned}$$

By [38], for each integer $c \geq 3$, there exists a restarting 1PFA \mathcal{P}_c recognizing EQ such that, in one pass along each given input $a^m b^n$, the probability of giving an erroneous decision (rejecting if the input should be accepted or accepting if it should be rejected) is at least c times smaller than the probability of giving a correct decision. From this moment on, one pass of \mathcal{P}_c along the given input $a^m b^n$ will be called a “round”. To be more exact, let \mathbb{A} and \mathbb{R} denote, respectively, the probability that \mathcal{P}_c accepts and rejects the input $a^m b^n$ in a single round. Then, by [38], the following holds:

$$\begin{aligned} \mathbb{R} &\leq \frac{1}{c} \cdot \mathbb{A}, & \text{if } m = n, \\ \mathbb{A} &\leq \frac{1}{c} \cdot \mathbb{R}, & \text{if } m \neq n. \end{aligned} \tag{1}$$

The remaining probability $\mathbb{N} = 1 - \mathbb{A} - \mathbb{R}$ represents computations ending with a “don’t know yet” result, after which \mathcal{P}_c restarts another round. In addition, we know from [38] that

$$\mathbb{A} = \frac{1}{3 \cdot (2c^2)^{m+n}}, \quad \text{independently of whether } m = n. \tag{2}$$

Note also that by combining (2) and (1) we get $\mathbb{A} > 0$ for $m = n$ and $\mathbb{R} > 0$ for $m \neq n$.

Finally, one of nice properties of these restarting 1PFAs is that they all use the same number of states, which is a constant that does not depend on c . (This is achieved by using arbitrarily small transition probabilities.)

Now we are ready to construct a 1PFA \mathcal{P}'_c for solving $\text{ExpEQ}(c)$. For the given input $(a^m b^n)^t$, where

$$t = 3 \cdot (2c^2)^{m+n} \cdot \lceil \ln c \rceil, \tag{3}$$

\mathcal{P}'_c simulates \mathcal{P}_c on $a^m b^n$ but, each time \mathcal{P}_c restarts its computation from the very beginning of $a^m b^n$, the machine \mathcal{P}'_c proceeds to the next copy of $a^m b^n$ along its own input. More precisely, if \mathcal{P}_c gives the decision of “acceptance” in the course of one round, \mathcal{P}'_c immediately accepts, by switching to an accepting state s_A . Similarly, if \mathcal{P}_c gives the decision of “rejection”, \mathcal{P}'_c immediately rejects, switching to a rejecting s_R . (In both these states, \mathcal{P}'_c consumes the rest of

the input by executing single-state loops.) Otherwise, \mathcal{P}'_c gets to the end of the current copy of $a^m b^n$ in a state s_N , representing here a “don’t know yet” outcome. That is, \mathcal{P}'_c is ready to retry the entire procedure with the next copy of $a^m b^n$. Clearly, along the input $(a^m b^n)^t$, this can be repeated t times, after which \mathcal{P}'_c halts in the state s_N .

Since the number of states in \mathcal{P}_c is a constant not depending on c , it is obvious that \mathcal{P}'_c also uses a constant number of states that does not depend on c .

Now, let \mathbb{A}_t and \mathbb{R}_t denote the respective probabilities that \mathcal{P}'_c accepts and rejects the input $(a^m b^n)^t$, by halting in the respective state s_A or s_R . The remaining probability $\mathbb{N}_t = 1 - \mathbb{A}_t - \mathbb{R}_t$ represents computations not deciding about acceptance or rejection, halting in the state s_N .

Let us estimate \mathbb{N}_t first. By combining $\mathbb{R} \geq 0$ with (2), (3), and the fact [20, Lm. A.3.60+] that $(1 - \frac{1}{x})^x < \frac{1}{e}$ for each real $x > 1$ (in that order), we get, independently of whether $m = n$, that

$$\begin{aligned} \mathbb{N}_t &= \mathbb{N}^t = (1 - \mathbb{A} - \mathbb{R})^t \leq (1 - \mathbb{A})^t = \left(1 - \frac{1}{3 \cdot (2c^2)^{m+n}}\right)^t = \left(1 - \frac{1}{3 \cdot (2c^2)^{m+n}}\right)^{3 \cdot (2c^2)^{m+n} \cdot \lceil \ln c \rceil} \\ &< \left(\frac{1}{e}\right)^{\lceil \ln c \rceil} \leq \frac{1}{c}. \end{aligned}$$

Therefore, using (1) and $\mathbb{A} > 0$, we can derive that the input $(a^m b^n)^t$ satisfying $m = n$ is accepted with probability at least

$$\begin{aligned} \mathbb{A}_t &= \sum_{i=1}^t \mathbb{A} \cdot \mathbb{N}^{i-1} = \mathbb{A} \cdot \frac{1 - \mathbb{N}^t}{1 - \mathbb{N}} = \frac{\mathbb{A}}{\mathbb{A} + \mathbb{R}} \cdot (1 - \mathbb{N}^t) \geq \frac{\mathbb{A}}{\mathbb{A} + \mathbb{A}/c} \cdot (1 - \mathbb{N}^t) = \frac{c}{c+1} \cdot (1 - \mathbb{N}^t) \\ &> \frac{c}{c+1} \cdot \left(1 - \frac{1}{c}\right) = 1 - \frac{2}{c+1}. \end{aligned}$$

Similarly, using (1) and $\mathbb{R} > 0$, the input $(a^m b^n)^t$ satisfying $m \neq n$ is rejected with probability at least

$$\begin{aligned} \mathbb{R}_t &= \sum_{i=1}^t \mathbb{R} \cdot \mathbb{N}^{i-1} = \mathbb{R} \cdot \frac{1 - \mathbb{N}^t}{1 - \mathbb{N}} = \frac{\mathbb{R}}{\mathbb{A} + \mathbb{R}} \cdot (1 - \mathbb{N}^t) \geq \frac{\mathbb{R}}{\mathbb{R}/c + \mathbb{R}} \cdot (1 - \mathbb{N}^t) = \frac{c}{c+1} \cdot (1 - \mathbb{N}^t) \\ &> \frac{c}{c+1} \cdot \left(1 - \frac{1}{c}\right) = 1 - \frac{2}{c+1}. \end{aligned}$$

Summing up, the success probability is always above $1 - \frac{2}{c+1}$ and consequently the error probability is always below $\frac{2}{c+1}$. (It should be pointed out that the standard PFAs do not use “don’t know” states. However, we can declare, by definition, the state s_N to be a rejecting state. This may potentially change all “don’t know” answers to errors, rejecting inputs that should be accepted. However, this does not change the fact that $\mathbb{A}_t > 1 - \frac{2}{c+1}$ for $m = n$, nor does it decrease \mathbb{R}_t for $m \neq n$.)

By taking $c = \max\{3, \lceil \frac{2}{\varepsilon} \rceil - 1\}$ for arbitrarily small but fixed $\varepsilon > 0$, we obtain the error probability $\frac{2}{c+1} \leq \varepsilon$, keeping the same constant number of states for each ε . This gives:

Theorem 6.3 *For each fixed $\varepsilon > 0$, there exists a promise problem solvable by a 1PFA with bounded error ε , using a constant number of states that does not depend on ε , but there is no 1DFA (hence, no other machine capable of recognizing only regular languages) solving the same problem.*

Proof: We only need to show that $\text{ExpEQ}(c)$ cannot be solved by any 1DFA, for no $c \geq 3$. For contradiction, let \mathcal{D} be a 1DFA solving $\text{ExpEQ}(c)$, for some $c \geq 3$. Without loss of generality, we assume that \mathcal{D} does not have undefined transitions, and hence it always halts at the end of the input. (Otherwise, we can define all missing transitions by switching to a single new rejecting state, in which \mathcal{D} scans the rest of the input.) Let $n \geq 1$ denote the number of states in \mathcal{D} .

Consider now the unary string a^n . Using the Claim presented in the proof of Theorem 6.1, we see that if \mathcal{D} traverses across a^n , starting in some state q_1 and ending in some q_2 , it will do the same also on the string $a^{n+h \cdot n!}$, for each $h \geq 1$. Clearly, the same holds for traversals of b^n and $b^{n+h \cdot n!}$.

Next, consider the input $(a^n b^n)^t$, where $t = 3 \cdot (2c^2)^{2n+2n!} \cdot \lceil \ln c \rceil$. If, on this input, \mathcal{D} halts in a state q' , then, by the observation above, \mathcal{D} must halt in the same state q' also on the inputs $(a^{n+n!} b^{n+n!})^t$ and $(a^n b^{n+2n!})^t$. Therefore, \mathcal{D} accepts $(a^{n+n!} b^{n+n!})^t$ if and only if it accepts $(a^n b^{n+2n!})^t$. But this is a contradiction, since the former string should be accepted by \mathcal{D} while the latter should be rejected. Consequently, there is no 1DFA solving $\text{ExpEQ}(c)$.

7 Final Remarks

A thorough study of promise problems can reveal several interesting properties of computational models and give new fundamental insights about them. In automata theory, promise problems have mainly been used to show how quantum models can do much better than the classical ones when compared to the case of language recognition. In this paper, we initiated a systematic work on promise problems for classical one-way finite automata (deterministic, nondeterministic, alternating, and probabilistic). In this context, we have also shown that randomness can do much better than other classical resources.

Promise problems can further be investigated for different computational models and from different perspectives. Two-way finite state machines and counter or pushdown automata models are the first ones coming to the mind. Moreover, we believe that some long-standing open problems, formulated for language recognition, might be solved more easily in the case of promise problems.

Acknowledgements. The authors thank Beatrice Palano for providing us a copy of [3].

References

- [1] Farid Ablayev, Aida Gainutdinova, Kamil Khadiev, and Abuzer Yakaryilmaz. Very narrow quantum OBDDs and width hierarchies for classical OBDDs. In *Descriptivity of Formal Systems*, volume 8614 of *Lecture Notes in Computer Science*, pages 53–64. Springer-Verlag, 2014.
- [2] Andris Ambainis and Abuzer Yakaryilmaz. Superiority of exact quantum automata for promise problems. *Information Processing Letters*, 112(7):289–291, 2012.
- [3] Maria Paola Bianchi, Carlo Mereghetti, and Beatrice Palano. Complexity of promise problems on classical and quantum automata. To appear.
- [4] R. G. Bukharaev. Probabilistic automata. *Journal of Mathematical Sciences*, 13(3):359–386, 1980.
- [5] Ashok K. Chandra, Dexter C. Kozen, and Larry J. Stockmeyer. Alternation. *Journal of the ACM*, 28(1):114–133, 1981.
- [6] Richard Chang, Juris Hartmanis, and Desh Ranjan. Space bounded computations: Review and new separation results. *Theoretical Computer Science*, 80:289–302, 1991.
- [7] Anne Condon and Richard J. Lipton. On the complexity of space bounded interactive proofs (extended abstract). In *FOCS'89: Proceedings of the 30th Annual Symposium on Foundations of Computer Science*, pages 462–467, 1989.
- [8] Pavol Ďuriš, Juraj Hromkovič, José D.P. Rolim, and Georg Schnitger. Las Vegas versus determinism for one-way communication complexity, finite automata, and polynomial-time computations. In *Symposium on Theoretical Aspects of Computer Science*, volume 1200 of *Lecture Notes in Computer Science*, pages 117–128. Springer-Verlag, 1997.

- [9] Cynthia Dwork and Larry Stockmeyer. Finite state verifiers I: The power of interaction. *Journal of the ACM*, 39(4):800–828, 1992.
- [10] Cynthia Dwork and Larry J. Stockmeyer. A time complexity gap for two-way probabilistic finite-state automata. *SIAM Journal on Computing*, 19(6):1011–1123, 1990.
- [11] Rūsiņš Freivalds. Probabilistic two-way machines. In *Mathematical Foundations of Computer Science*, volume 118 of *Lecture Notes in Computer Science*, pages 33–45. Springer-Verlag, 1981.
- [12] Viliam Geffert. Nondeterministic computations in sublogarithmic space and space constructibility. *SIAM Journal on Computing*, 20:484–498, 1991.
- [13] Viliam Geffert. An alternating hierarchy for finite automata. *Theoretical Computer Science*, 445:1–24, 2012.
- [14] Viliam Geffert and Alexander Okhotin. Transforming two-way alternating finite automata to one-way nondeterministic automata. In *Mathematical Foundations of Computer Science, Part I*, volume 8634 of *Lecture Notes in Computer Science*, pages 291–302. Springer-Verlag, 2014.
- [15] Viliam Geffert and Abuzer Yakaryılmaz. Classical automata on promise problems. In *Descriptional Complexity of Formal Systems*, volume 8614 of *Lecture Notes in Computer Science*, pages 126–137. Springer-Verlag, 2014.
- [16] Oded Goldreich. On promise problems: A survey. In *Essays in Memory of Shimon Even*, volume 3895 of *Lecture Notes in Computer Science*, pages 254–290. Springer-Verlag, 2006.
- [17] Jozef Gruska, Daowen Qiu, and Shenggen Zheng. Generalizations of the distributed Deutsch-Jozsa promise problem. Technical report, 2014. arXiv:1402.7254.
- [18] Jozef Gruska, Daowen Qiu, and Shenggen Zheng. Potential of quantum finite automata with exact acceptance. Technical Report arXiv:1404.1689, 2014.
- [19] John E. Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Prentice Hall, 2007.
- [20] Juraĳ Hromkoviĉ. *Design and Analysis of Randomized Algorithms (Introduction to Design Paradigms)*. Springer-Verlag, 2005.
- [21] Juraĳ Hromkoviĉ and Georg Schnitger. On the power of Las Vegas for one-way communication complexity, OBDDs, and finite automata. *Information and Computation*, 169(2):284–296, 2001.
- [22] Galina Jirásková and Giovanni Pighizzini. Optimal simulation of self-verifying automata by deterministic automata. *Information and Computation*, 209:528–535, 2011.
- [23] Christos A. Kapoutsis. Removing bidirectionality from nondeterministic finite automata. In *Mathematical Foundations of Computer Science*, volume 3618 of *Lecture Notes in Computer Science*, pages 544–555. Springer-Verlag, 2005.
- [24] Christos A. Kapoutsis, Richard Kráľoviĉ, and Tobias Mömke. Size complexity of rotating and sweeping automata. *Journal of Computer and System Sciences*, 78(2):537–558, 2012.
- [25] Yu. I. Kuklin. Two-way probabilistic automata. *Avtomatika i vychislitel'naja tekhnika*, 5:35–36, 1973. (Russian).

- [26] Rajendra Kumar. *Theory of Automata, Languages, and Computation*. Tata McGraw-Hill, 2010.
- [27] Richard E. Ladner, Richard J. Lipton, and Larry J. Stockmeyer. Alternating pushdown and stack automata. *SIAM Journal on Computing*, 13(1):135–155, 1984.
- [28] Carlo Mereghetti and Giovanni Pighizzini. Optimal simulations between unary automata. *SIAM Journal on Computing*, 30:1976–1992, 2001.
- [29] Yumiko Murakami, Masaki Nakanishi, Shigeru Yamashita, and Katsumasa Watanabe. Quantum versus classical pushdown automata in exact computation. *IPSJ Digital Courier*, 1:426–435, 2005.
- [30] Masaki Nakanishi. Quantum pushdown automata with a garbage tape. Technical Report arXiv:1402.3449, 2014.
- [31] Azaria Paz. *Introduction to Probabilistic Automata*. Academic Press, New York, 1971.
- [32] Michael O. Rabin. Probabilistic automata. *Information and Control*, 6:230–243, 1963.
- [33] Jibran Rashid and Abuzer Yakaryilmaz. Implications of quantum automata for contextuality. In *Conference on Implementation and Application of Automata*, volume 8587 of *Lecture Notes in Computer Science*, pages 318–331. Springer-Verlag, 2014. arXiv:1404.2761.
- [34] Klaus Reinhardt and Abuzer Yakaryilmaz. The minimum amount of useful space: New results and new directions. In *Developments in Language Theory*, volume 8633 of *Lecture Notes in Computer Science*, pages 315–326. Springer-Verlag, 2014.
- [35] Michael Sipser. Lower bounds on the size of sweeping automata. *Journal of Computer and System Sciences*, 21(2):195–202, 1980.
- [36] Richard Edwin Stearns, Juris Hartmanis, and Philip M. Lewis II. Hierarchies of memory limited computations. In *IEEE Conference Record on Switching Circuit Theory and Logical Design*, pages 179–190, 1965.
- [37] John Watrous. Quantum computational complexity. In Robert A. Meyers, editor, *Encyclopedia of Complexity and Systems Science*, pages 7174–7201. Springer-Verlag, 2009.
- [38] Abuzer Yakaryilmaz and A. C. Cem Say. Succinctness of two-way probabilistic and quantum finite automata. *Discrete Mathematics and Theoretical Computer Science*, 12(2):19–40, 2010.
- [39] Shenggen Zheng, Jozef Gruska, and Daowen Qiu. On the state complexity of semi-quantum finite automata. In *Language and Automata Theory and Applications*, volume 8370 of *Lecture Notes in Computer Science*, pages 601–612. Springer-Verlag, 2014.
- [40] Shenggen Zheng, Daowen Qiu, Jozef Gruska, Lvzhou Li, and Paulo Mateus. State succinctness of two-way finite automata with quantum and classical states. *Theoretical Computer Science*, 499:98–112, 2013.